

CLAIMS

What is claimed is:

1 1. A switch fabric for routing data from one or more sources towards one or more destinations,
2 comprising a switching stage configured between an input stage and an output stage, wherein:
3 the input stage is configured to receive the data transmitted from the one or more sources and
4 forward the data to the switching stage;
5 the switching stage is configured to route the data received from the input stage to the output stage;
6 the output stage is configured to transmit the data received from the switching stage towards the one
7 or more destinations;
8 the input stage has one or more input ports, wherein each input port can be configured to receive data
9 from a source;
10 the output stage has one or more output ports, wherein each output port can be configured to transmit
11 data towards a destination;
12 the input stage comprises a plurality of input routing queues configured to store the data until the data
13 is ready to be forwarded to the switching stage;
14 the output stage comprises a plurality of output routing queues configured to store the data until the
15 data is ready to be transmitted towards the one or more destinations;
16 at least one input port can be programmably configured to store data in two or more input routing
17 queues that are associated with a single output port; and
18 at least one output port can be programmably configured to receive data from two or more output
19 routing queues that are associated with a single input port.

1 2. The invention of claim 1, wherein:
2 different sets of data received at an input port from a single source can be stored in two or more
3 different input routing queues; and
4 different sets of data to be transmitted from an output port towards a single destination can be stored
5 in two or more different output routing queues.

1 3. The invention of claim 2, wherein the data received at the input port from the single source can
2 be separated into the different sets of data based on different destinations.

1 4. The invention of claim 2, wherein the data received at the input port from the single source can
2 be separated into the different sets of data based on different quality-of-service levels.

1 5. The invention of claim 2, wherein the data received at the input port from the single source can
2 be separated into the different sets of data to achieve isolation.

1 6. The invention of claim 1, wherein:

2 each input port can be programmably configured independent of each other input port; and
3 each output port can be programmably configured independent of each other output port.

1 7. The invention of claim 1, wherein:

2 the input stage comprises one or more input devices, wherein each input device comprises a plurality
3 of input ports and a plurality of input routing queues;

4 the switching stage comprises one or more crossbar devices;

5 the output stage comprises one or more output devices, wherein each output device comprises a
6 plurality of output ports and a plurality of output routing queues;

7 each input port of each input device can be programmably configured to store data in two or more
8 input routing queues of that input device; and

9 each output port of each output device can be programmably configured to receive data from two or
10 more output routing queues of that output device.

11 8. The invention of claim 7, wherein:

12 each input device can be configured independent of each other input device;

13 each output device can be configured independent of each other output device;

14 each input port of each input device can be programmably configured independent of each other
15 input port of that input device; and

16 each output port of each output device can be programmably configured independent of each other
17 output port of that output device.

1 9. The invention of claim 7, wherein:

2 a maximum number of the input routing queues in each input device that can be configured to receive
3 data from an input port of that input device is a function of (1) the number of input ports of that input
4 device that are configured to receive data from the one or more sources and (2) the total number of output
5 ports in the output stage that are configured to transmit data towards the one or more destinations; and

6 a maximum number of the output routing queues in each output device that can be configured to
7 transmit data to an output port of that output device is a function of (1) the number of output ports of that

8 output device that are configured to transmit data towards the one or more destinations and (2) the total
9 number of input ports in the input stage that are configured to receive data from the one or more sources.

1 10. The invention of claim 1, wherein the switch fabric is fabricated as one or more integrated
2 circuits.

1 11. A switch fabric for routing data from one or more sources towards one or more destinations,
2 comprising a switching stage configured between an input stage and an output stage wherein:
3 the input stage is configured to receive the data transmitted from the one or more sources and
4 forward the data to the switching stage;
5 the switching stage is configured to route the data received from the input stage to the output stage;
6 the output stage is configured to transmit the data received from the switching stage towards the one
7 or more destinations;
8 the input stage transmits bids to the switching stage to request connections through the switching
9 stage for routing the data;
10 the output stage is configured to transmit status information about the output stage to the input stage;
11 and
12 the input stage is configured to generate the bids transmitted to the switching fabric based on the
13 status information about the output stage.

1 12. The invention of claim 11, wherein:
2 the output stage is further configured to transmit status information about the output stage to the
3 switching stage; and
4 the switching stage is configured to determine whether to accept or reject the bids based on the status
5 information about the output stage.

1 13. The invention of claim 12, wherein:
2 the output stage comprises a plurality of output routing queues configured to store the data prior to
3 being transmitted towards the one or more destinations;
4 the output stage has a plurality of output ports, wherein each output port can be configured to
5 transmit data towards a destination;
6 the status information about the output stage that is used by the input stage corresponds to per queue
7 status information for each output routing queue; and

8 the status information about the output stage that is used by the switching stage corresponds to per
9 port status information for each output port.

1 14. The invention of claim 11, wherein:

2 the output stage comprises a plurality of output routing queues configured to store the data prior to
3 being transmitted towards the one or more destinations;

4 the output stage has a plurality of output ports, wherein each output port can be configured to
5 transmit data towards a destination;

6 the status information about the output stage that is used by the input stage corresponds to (1) per
7 queue status information for each output routing queue and (2) per port status information for each output
8 port.

1 15. The invention of claim 11, wherein:

2 the input stage comprises a plurality of input devices, each input device performing a port expansion
3 function;

4 the switching stage comprises one or more crossbar devices, each crossbar device performing a data
5 routing function;

6 the output stage comprises a plurality of output devices, each output device performing a port
7 contraction function;

8 the status information about each output device is transmitted to all of the input devices in the input
9 stage; and

10 each input device is configured to generate bids transmitted to the switching fabric based on the
11 status information about all of the output devices.

1 16. The invention of claim 15, wherein:

2 each output device comprises a plurality of output routing queues and a plurality of output ports;

3 the status information transmitted from each output device to all of the input devices comprises both
4 per queue status information for each output routing queue of that output device and per port status
5 information for each output port of that output device.

1 17. The invention of claim 16, wherein:

2 each input device avoids generating bids for output routing queues that are congested as indicated by
3 the per queue status information; and

4 each input device avoids generating bids for output ports that are congested as indicated by the per
5 port status information.

1 18. The invention of claim 11, wherein the switch fabric is fabricated as one or more integrated
2 circuits.

1 19. A switch fabric for routing data from one or more sources towards one or more destinations,
2 comprising a switching stage configured between an input stage and an output stage wherein:

3 the input stage is configured to receive the data transmitted from the one or more sources and
4 forward the data to the switching stage;

5 the switching stage is configured to route the data received from the input stage to the output stage;

6 the output stage is configured to transmit the data received from the switching stage towards the one
7 or more destinations;

8 the input stage is configured to transmit bids to the switching stage to request connections through
9 the switching stage for routing the data to the output stage;

10 the switching stage is configured to determine whether to accept or reject each bid and to transmit a
11 grant/rejection signal to the input stage identifying whether the bid is accepted or rejected;

12 for each rejected bid, the grant/rejection signal explicitly identifies a reason for rejecting the bid; and

the input stage determines how to react to a rejected bid based on the reason the bid was rejected.

1 20. The invention of claim 19, wherein:

2 the input stage comprises a plurality of input devices, each input device performing a port expansion
3 function;

4 the switching stage comprises one or more crossbar devices, each crossbar device performing a data
5 routing function;

6 the output stage comprises a plurality of output devices, each output device performing a port
7 contraction function;

8 each input device is configured to transmit bids to the one or more crossbar devices to request
9 connections through the switching stage for routing the data to the output devices; and

10 each crossbar device is configured to determine whether to accept or reject each bid received from an
11 input device and the transmit a grant/rejection signal to that input device identifying whether the bid is
12 accepted or rejected.

1 21. The invention of claim 20, wherein each input device is configured:

2 to transmit, if an initial bid is rejected due to bid collision, a subsequent bid for the same output
3 device; and
4 to transmit, if an initial bid is rejected due to back-pressure, a subsequent bid for a different output
5 device.

1 22. The invention of claim 21, wherein the back-pressure comprises both per queue status
2 information for each output routing queue in each output device and per port status information for each
3 output port in each output device.

1 23. The invention of claim 20, wherein, in each input device, bids for each output device are
2 organized and ordered by traffic class.

1 24. The invention of claim 20, wherein each input device is configured to reschedule a low-priority
2 cell if a bid for a corresponding routing queue is rejected.

1 25. The invention of claim 19, wherein the switch fabric is fabricated as one or more integrated
2 circuits.